IN THE SPECIFICATION

Please amend the Specification as follows:

(1) Amend the text beginning at page 4, line 25 and ending at page 5, line 2 as follows:

-- The present invention is a methodology for providing fault detection and service restoration for a multiservice switch on a per flow basis. Referring still to Fig. 1, ingress data (e.g., communicating through the first service shelf 12 and one of the on-line service cards 24) transmits the same data over each of the redundant cores 16, 18. An egress receiver 25 (in the second service shelf ~~13~~ 20) selects on a per path basis which core to utilize. Bi-directional paths are not necessarily grouped together. That is, for a duplex path, one direction of transmission can proceed through the first core 16 and the other direction can proceed through the other core 18, if required. --

(2) Amend the text beginning at page 6, line 10 and ending at page 6, line 7 as follows:

-- The aggregator device 38 acts as an interface between the service cards ~~12~~ 24, 26 and the switching core and essentially distributes core traffic throughout the service shelf. The aggregator 38 acts as a datapath flow switch, directing flows to either the normally active service card slot or to the dedicated protection slot. Note that neither core bandwidth, nor bandwidth of the service cards (shown in greater detail in Fig. 6) is wasted by the aggregator cross-connect function to the service cards (Fig. 6). In all cases, the aggregator 38 will allow control information connectivity through the core to all attached service cards 24, 26 and shelf control processors 36. Although shown and described as an applications specific integrated circuit (ASIC), it would be understood that the functionality of the aggregator 38 as described herein may also be implemented using discrete components. As shown in Fig. 2 and 3, the core side of the aggregator 38 couples to multiple serializer/deserializer blocks 40. The implementation and function of a serializer/deserializer would be well known to a person skilled in the art. The serializer/deserializers 40 couple to

optical/electrical (O/E) components 42 in order to provide the interface to the switching core. Failure of a link will be detected by a serializer/deserializer 40 or the aggregator device 38 and reported to the shelf control processor 36 through a control interface on the aggregator. Failures may be detected, for example, by the loss of a clock signal corresponding to the link or an invalid parity across the link. Other types of failures that are detectable and that can be characterized as a link failure would be apparent to those skilled in the art. As will be explained, the shelf control processor 36 (in combination with the aggregator 38) trigger appropriate corrective action in response to a link failure. The aggregator 38 on the core interface card 14 also contains a thread switch function 44 for service card protection. The switch function 44 allows the core interface card 14 to steer traffic on a given thread to/from an active service card to a protection card. For the shelf, service card protection will be 1:2. The core interface card 14 (and the shelf control processor 36) will control the protection switching of the interface. In addition, as will be explained, an arbiter function on the service card can detect link failures on the basis, for example, of the receipt/non-receipt of link test cells. --

(3) Amend the text at page 7, lines 8-22 inclusive as follows:

-- Fig. 4 shows a functional block diagram of the aggregator device 38. The aggregator 38 includes ingress receive logic 50 and egress transmit logic 52 on the service card side. ~~Ingress~~ Egress transmit logic 54 and egress receive logic 56 are also found on the core side of the aggregator 38. There are two aggregation functions – AGR0 and AGR1 – implemented in the aggregator (AGR) ASIC, each performing an aggregation of up to 6 independent data streams into, for example, a 2.5Gbps or higher thread. These two aggregation functions are independent and the operation of one does not affect any state of the other. In one exemplary embodiment, each aggregator function AGR0, AGR1 includes a multiplexer unit 58 which couples to the ingress receive logic 50, a cell decode unit 60 which couples to the output of the multiplexer 58 and a

buffer management unit 62 which couples to the output of the cell decode unit 60. A credit/grant manager function 64 and a multicast unit 66 each couple to the output of the buffer management unit 62. A virtual output queue (VOQ) memory interface 68 and a pointer memory interface 70 each couple to the multicast unit 66. A VOQ scheduler 72 couples to the credit/grant manager 64.
--

(4) Amend the text beginning at page 10, line 29 and ending at page 11, line 10 as follows:

-- Referring to Fig. 7, a functional block diagram of the ARB ASIC 76 is shown. The exemplary embodiment of the ARB includes six interfaces: a PCI (processor interface) interface, a physical layer interface (PI Sched RX and TX), a SAR interface (RX and TX), two AGR interfaces (RX and TX, one per core) and an external memory interface. As discussed previously, the ARB includes a link test cell generator 102 and a link test cell receiver 104 which will be used in the system to verify flow integrity. The link test cell (LTC) generator 102 and receiver 104 couple to the aggregator interface 106, the link test cell receiver 104 coupling through respective egress filters 108. The ARB also includes internal priority queues (four QOS levels) 110 for egress traffic, the inputs of which couple to the egress filter 108. The priority queues couple to egress transmit ports (TM and Utopia) 112, 114 through a scheduler ~~116 or~~ 118. The egress filters 108 in the ARB provide a filtering function that is used to determine if the ARB should accept unicast and multicast cells from the AGRs.